arXiv:1409.0742v3 [cs.CC] 10 Aug 2015

New Algorithms and Hard Instances for Non-Commutative Computation

Christian Engels* B. V. Raghavendra Rao†

August 11, 2015

Recent developments on the complexity of the non-commutative determinant and permanent [Chien et al. STOC 2011, Bläser ICALP 2013, Gentry CCC 2014] have settled the complexity of non-commutative determinant with respect to the structure of the underlying algebra. Continuing the research further, we look to obtain more insights on hard instances of non-commutative permanent and determinant.

We show that any Algebraic Branching Program (ABP) computing the Cayley permanent of a collection of disjoint directed two-cycles with distinct variables as edge labels requires exponential size. For graphs where every connected component contains at most six vertices, we show that evaluating the Cayley permanent over any algebra containing 2×2 matrices is #P complete.

Further, we obtain efficient algorithms for computing the Cayley permanent/determinant on graphs with bounded component size, when vertices within each component are not far apart from each other in the Cayley ordering. This gives a tight upper and lower bound for size of ABPs computing the permanent of disjoint two-cycles. Finally, we exhibit more families of non-commutative polynomial evaluation problems that are complete for #P.

Our results demonstrate that apart from the structure of underlying algebras, relative ordering of the variables plays a crucial role in determining the complexity of non-commutative polynomials.

1. Introduction

Background

The study of algebraic complexity theory was initiated by Valiant in his seminal paper [23] where he showed that computing the permanent of an integer matrix is #P complete.

*Saarland University, Department of Computer Science, Germany engels@cs.uni-saarland.de
†IIT Madras, Department of Computer Science and Engineering, Chennai, India, bvrr@cse.iitm.ac.in

Since then, separating the complexities of permanent and determinant has been the focal point of this research area which led to the development of several interesting results and techniques. (See [8, 22] for good surveys on these topics.)

The underlying ring plays an important role in algebraic complexity theory. While the research focused mainly on the permanent vs determinant problem over fields and commutative rings there has also been an increasing amount of interest over non-commutative algebras. Nisan [19] was the first to consider the complexity of these two polynomials over non-commutative algebras. He showed that any non-commutative arithmetic formula over the free $\mathbb{K}$ algebra computing the permanent or determinant of an $n \times n$ matrix requires size $2^{\Omega(n)}$ where $\mathbb{K}$ is any field. Later on, this was generalized to other classes of algebras in [9]. More recently, Limaye, Malod and Srinivasan [16] generalized Nisan's technique to prove lower bounds against more general classes of non-commutative circuits. Nisan's work left the problem of determining the arithmetic circuit complexity of non-commutative determinant as an open question.

In a significant breakthrough, Arvind and Srinivasan [4] showed that computing the Cayley determinant is #P hard over certain matrix algebras.Finally this question was settled by Bläser [7] who classified such algebras. Further, Gentry [14] simplified these reductions.

Motivation

Though the studies in [4, 7] highlight the role of the underlying algebra in determining the complexity of the non-commutative determinant they do not shed much light on the combinatorial structure of non-commutative polynomials that are #P hard. One could ask: *Does the hardness stem from the underlying algebra or are there inherent properties of polynomials that make them #P hard in the non-commutative setting?* Our results in this paper indicate that relative ordering among the variable also plays an important role in the hardness of certain non-commutative polynomials.

As a first step, we look for polynomials that are easier to compute than the determinant in the commutative setting and whose non-commutative versions are #P hard. Natural candidate polynomials are the elementary symmetric polynomials and special cases of determinant/permanent. One way to obtain special cases of determinant/permanent would be to restrict the structure of the underlying graph. For example, let G be a directed graph consisting of n cycles $(0,1),(2,3),\ldots,(2n-2,2n-1)$ of length two with self loops where each edge is labeled by a distinct variable. The permanent of G, $\mathsf{perm}(G)$, is given by $\prod_{i=0}^{n-1}(x_{2i,2i}x_{2i+1,2i+1} + x_{2i,2i+1}x_{2i+1,2i})$ where $x_{i,j}$ is the variable labeling of the edge (i,j). This is one of the easiest to compute but non trivial special case of permanent.

1.0.1. Our Results

We study the complexity of the Cayley permanent (C-perm) on special classes of graphs. We exhibit a family of collections of disjoint two-cycles for which any algebraic branching program (ABP) computing the C-perm should have size $2^{\Omega(n)}$ (Corollary 4.1). Further, we

exhibit a parameter $\mathsf{cut}(G)$ (see Section 4 for the definition) for a collection G of disjoint two-cycles on n vertices such that any ABP computing $\mathsf{C\text{-}perm}(G)$ has size $2^{\Theta(\mathsf{cut}(G))}$ (Theorem 4.2). This makes the lower bound in Corollary 4.1 tight up to a constant factor in the exponent. It should be noted that our results also hold for the case of the Cayley determinant ($\mathsf{C\text{-}det}$) on such graphs. We also observe that for graphs of component size greater or equal to six the problem of evaluating $\mathsf{C\text{-}perm}$ is #P complete (Theorem 5.1).

On the positive side, for graphs where each strongly connected component has at most c vertices we obtain an ABP of size $n^{O(c)}c^{\mathsf{near}(G)}$ computing the $\mathsf{C\text{-}perm}$ (Theorem 3.1) where $\mathsf{near}(G)$ is a parameter (see Definition 3.1) depending on the labeling of vertices on the graph.

We demonstrate a non-commutative variant of the elementary symmetric polynomial that is #P hard over certain algebras (Theorem 5.3). Finally, we show that computing $\mathsf{C\text{-}perm}$ on rank one matrices is #P hard.

1.0.2. Related Results

The study of commutative permanent on special classes of matrices was initiated by Barvinok [6] who gave a polynomial time algorithm for computing the permanent of rank one matrices over a field. More recently, Flarup, Koiran and Lyaudet [11] showed that computing the permanent of bounded tree-width graphs can be done by polynomial size formulas. This was further extended by Flarup and Lyaudet [12] to other width measures on graphs. Datta et al. [10] showed that computing the permanent on planar graphs is as hard as the general case.

1.0.3. Comparison to other results

Results reported in [4, 7, 14] highlight the importance of the underlying algebra and characterizes algebras for which $\mathsf{C\text{-}det}$ is #P hard. In contrast, our results shed light on the role played by the order in which vertices are labeled in a graph. For example, the commutative permanent of disjoint two-cycles has a depth three formula given by $\prod_{i=0}^{n-1}(x_{2i,2i}x_{2i+1,2i+1}+x_{2i,2i+1}x_{2i+1,2i})$ whereas $\mathsf{C\text{-}perm}$ on almost all orderings of vertices requires exponential size ABPs.

2. Preliminaries

For definitions of complexity classes the reader is referred to any of the standard text books on Computational Complexity Theory, e.g., [1]. Let $\mathbb{K}$ be a field and $S = \mathbb{K}[x_1, \ldots, x_n]$ be the ring of polynomials over $\mathbb{K}$ in n variables. Let R denote a non-commutative ring with identity and associativity property. Unless otherwise stated, we assume that R is an algebra over $\mathbb{K}$ and contains the algebra of $n \times n$ matrices with entries from $\mathbb{K}$ as a subalgebra.

An *arithmetic circuit* is a directed acyclic graph where every vertex has an in-degree either zero or two. Vertices of zero in-degree are called *input* gates and are labeled by elements in $R \cup \{x_1 \ldots, x_n\}$. Vertices of in-degree two are called *internal* gates and have

their labels from $\{\times, +\}$. An arithmetic circuit has at least one vertex of out degree zero called an *output* gate. We assume that an arithmetic circuit has exactly one output gate. A polynomial p_g in $R[x_1, \ldots, x_n]$ can be associated with every gate g of an arithmetic circuit defined in an inductive fashion. Input gates compute their label. Let g be an internal gate with left child f and right child h, then $p_g = p_f \ \mathsf{op} \ p_h$ where op is the label of g. The polynomial computed by the circuit is the polynomial at one of the output gates and denoted by p_C. The size of an arithmetic circuit is the number of gates in it and is denoted by $\mathsf{size}(C)$. We restrict ourselves to circuits where coefficients of the polynomials computed at every gate can be represented in at most $\mathsf{poly}(\mathsf{size}(C))$ bits.

An *algebraic branching program* (ABP) is a directed acyclic graph with two special nodes s, t and edges labeled by variables or constants in R. The weight of a path is the product of the weights of its edges. The polynomial computed by an ABP P is the sum of the weights of all $s \rightsquigarrow t$ paths in P, and is denoted by p_P.

Over a non-commutative ring, there are many possibilities for defining the determinant/permanent of a matrix depending on the ordering of the variables (see for example [5]). We will use the well known definitions of the *Cayley determinant* and *Cayley permanent*. Let $X = (x_{i,j})_{1 \le i,j \le n}$ be an $n \times n$ matrix with distinct variables $x_{i,j}$. Then

$$\mathsf{C\text{-}det}(X) = \sum_{\sigma \in S_n} \mathrm{sgn}(\sigma) x_{1,\sigma(1)} \cdots x_{n,\sigma(n)}; \quad \text{and } \mathsf{C\text{-}perm}(X) = \sum_{\sigma \in S_n} x_{1,\sigma(1)} \cdots x_{n,\sigma(n)}.$$

In the above, S_n denotes the set of all permutations on n symbols. Note that $\mathsf{C\text{-}det}$ and $\mathsf{C\text{-}perm}$ can also be seen as functions taking $n \times n$ matrices with entries from R as input. Given a weighted directed graph G on n vertices with weight $x_{i,j}$ for the edge $(i,j) \in E(G)$, the Cayley permanent of G denoted by $\mathsf{C\text{-}perm}(G)$ is the permanent of the weighted adjacency matrix of G. It is known that [8] $\mathsf{C\text{-}perm}(G)$ is the sum of the Cayley weights of all cycle covers of G.

The tensor product of two matrices $A, B \in \mathbb{K}^{n \times n}$ with entries $a_{i,j}, b_{i,j}$ is denoted by $A \otimes B$ and is given by

$$A \otimes B = \begin{pmatrix} a_{1,1}B & \cdots & a_{1,n}B \\ \vdots & \ddots & \vdots \\ a_{n,1}B & \cdots & a_{n,n}B \end{pmatrix}.$$

Let P be an ABP over disjoint sets of variables $X \cup Y$, with $|X| = n$ and $|Y| = m$. Let $p_P(X, Y)$ be the polynomial computed by P. P is said to be *read once certified* [17] in Y if there are numbers $0 = i_0 < i_1 < \cdots < i_m$ where i_m is at most the length of P and there is a permutation $\pi \in S_m$ such that between layers from i_j to i_{j+1} no variable other than $y_{\pi(j+1)}$ from the set Y appears as a label. We use the following result from [17]. The proof given in [17] works only in the commutative setting, see Appendix B for the non-commutative case.

Proposition 1 ([17]). *Let P be an ABP on $X \cup Y$ read-once certified in Y. Then the polynomial $\sum_{e_1, e_2, \ldots, e_m \in \{0,1\}^m} p_P(X, e_1, \ldots, e_m)$ can be computed by an ABP of size $\mathsf{poly}(\mathsf{size}(P))$.*

Let $\mathcal{A}$ be a non-deterministic s-space bounded algorithm that uses non-deterministic bits in a read-once fashion and outputs a monomial on each of the accepting paths. We assume that a non-commutative monomial is output as a string in a write-only tape and non-deterministic paths are represented by binary strings $e \in \{0,1\}^m$, $m \leq 2^{O(s)}$. The polynomial $p_{\mathcal{A}}$ computed by $\mathcal{A}$ is the sum of the monomial output on each of the accepting paths of $\mathcal{A}$, i.e., $p(x_1, \ldots, x_n) = \sum_e A(x_1, \ldots, x_n, e)$, where the sum is taken over all accepting paths e of $\mathcal{A}$, and $A(x_1, \ldots, x_n, e)$ denotes the monomial output along path represented by e.

Proposition 2 (folklore)**.** *Let $A(X)$ be an s-space bounded non-deterministic algorithm as above. There is a non-commutative ABP P of size $2^{O(s)}$ that computes the polynomial $p_A(X)$.*

3. An algorithm for Cayley Permanent

In this section, we give an algorithm for C-perm that is parameterized by the maximum difference between labelings of vertices in individual components.

In what follows, we identify the vertices of a graph with the set $[n]$. A directed graph G on n vertices is said to have *component size bounded by c* if every strongly connected component of G contains at most c vertices where $c > 0$. We assume that edges of G are labeled by distinct variables. Firstly, we define a parameter that measures the closeness of labelings in each component.

Definition 3.1. *Let G be a directed graph. The* nearness *parameter* $\mathsf{near}(C)$ *of a strongly connected component C of G is defined as* $\mathsf{near}(C) = \max_{i,j \in C} |i - j|$*. The nearness parameter of G is defined as* $\mathsf{near}(G) = \max_C \mathsf{near}(C)$*, where the maximum is taken over the set of all strongly connected components in G.*

Theorem 3.1. *Let G be a directed graph with component size bounded by c and edges labeled by distinct variables. Then there exists an ABP of size $n^{O(c)} c^{\mathsf{near}(G)}$ computing the Cayley permanent of the adjacency matrix of G.*

Proof. For an edge $(i,j) \in E(G)$, let $x_{i,j}$ denote the variable label on (i,j). Let A_G be the weighted adjacency matrix of G. Note that, the Cayley permanent of A_G equals the sum of weights of cycle covers in G where the weight of a cycle cover γ is the product of labels of edges in γ multiplied in the Cayley order.

We describe a non-deterministic small-space bounded procedure P that guesses a cycle cover γ in G and outputs the product of weights of γ with respect to the Cayley ordering as a string of variables. Additionally, we ensure that the algorithm P uses the non-deterministic bits in a read-once fashion, and by the closure property of ABP under read-once exponential sums (c.f. Proposition 1), we obtain the required ABP. Suppose $C_1, \ldots, C_r$ are the strongly connected components of G, sorted in the ascending order of the smallest vertex in each component. Then any cycle cover γ of G can be decomposed into cycle cover γ_i of the component C_i. The only difficulty in computing the weight of γ is the Cayley ordering of the variables. However, with a careful implementation,

we show that this can be done in space $O(\log c \cdot \mathsf{near}(G) + \log n)$. We represent a cycle cover in G as a permutation γ where $\gamma(i)$ is the successor of vertex i in the cycle cover represented by γ. We begin with the description of the non-deterministic procedure P. Let T represent the set of vertices v in the partial cover that is being built by the procedure where the weight of the edge going out of v is not yet output, and pos the current position going from 1 to n.

1. Initialize pos $:= 1$, $T := \emptyset$, $\gamma :=$ the cycle cover of the empty graph, $f = 1$.
2. For $1 \leq i \leq r$ repeat steps 3 & 4.
3. Non-deterministically guess a cycle cover γ' in C_i, and set $\gamma = \gamma \uplus \gamma'$, $T = T \cup V(C_i)$ where $V(C_i)$ is the set vertices in C_i.
4. While there is a vertex $k \in T$ with $k =$ pos do the following:

   Set $f = f \cdot x_{k,\gamma(k)}$; pos $:=$ pos $+ 1$; and $T := T \setminus \{k\}$.
5. If pos $= n$, then output f and accept.

Let $\mathsf{Acc}(G)$ be the sum of the monomials output by the algorithm on all accepting paths.

Claim 1. $\mathsf{Acc}(G) = \mathsf{C\text{-}perm}(G)$. *Moreover, the algorithm P uses $O(\log c \cdot \mathsf{near}(G) + \log n)$ space, and is read-once on the non-deterministic bits.*

of the Claim. Recall that a permutation $\gamma \in S_n$ is a cycle cover of G if and only if it can be decomposed into vertex disjoint cycle covers $\gamma_1, \ldots, \gamma_r$ of the strongly connected components $C_1, \ldots, C_r$ in G. Thus Step 3 enumerates all possible cycle covers in G. Also, the weights output at every accepting path are in the Cayley order.

We have $T = \{k \mid \text{pos} < k \text{ and } k \text{ occurs in the components already explored}\}$. Firstly, we argue that at any point in time in the algorithm, $|T| \leq \mathsf{near}(G) + c$. Suppose the algorithm has processed components up to C_i and is yet to process C_{i+1}. Let $\mu = \max_{v \in T} v$. Since the components are in ascending order with respect to the smallest vertex in them, the component C_j with $\mu \in C_j$ must have $\mathsf{near}(C_j) \geq \mu - \text{pos}$. Thus $\mu - \text{pos} \leq \mathsf{near}(G)$. Also, just before step 3 in any iteration, for any $v \in T$, we have pos $< v \leq \mu$ and hence $|T| \leq \mu - \text{pos} + c \leq \mathsf{near}(G) + c$.

Note that it is enough to store the labels of the vertices in T and the choice $\gamma(v)$ made during the non-deterministic guess for each $v \in T$ and hence $O(|T| \log n)$ additional bits of information needs to be stored. However, we will show that it is possible to implement the algorithm without explicitly remembering the vertices in T and using only $O(|T| \log c)$ additional bits in memory. Suppose that the vertices in T are ordered as they appear in $C_1, C_2, \ldots, C_r$ where vertices within a component are considered in the ascending order of their labels. Let B be a vector of length $\mathsf{near}(G)$ where each entry B_j is $\log c$ bits long which indicates the neighbour of the jth vertex in T. Now, we show how to implement step 4 in the procedure using B as a data structure for T. To check if there is a $k \in T$ with $k =$ pos, we can scan the components from $C_1, \ldots, C_i$ and check if the vertex assigned to pos occurs in one of the components. Remember that $\gamma(k)$ is the

successor of k in the cycle cover γ. To obtain $\gamma(k)$ from B, we need to know the number j of vertices v that appear in components $C_1, \ldots, C_i$ such that $v \geq \text{pos}$ and that occur before k. Then $\gamma(k) = B_j$. Once B_j is used, we remove B_j from B and shift the array $B_{j+1}, \ldots B_{\mathsf{near}(G)+c}$ by one index towards the left. Further, we can implement step 3 by simply appending the information for $V(C_i)$ given by γ' to the right of the array B. We require at most $O(c \log n)$ bits of space guessing a cycle cover γ_i for component C_i which can be re-used after the non-deterministic guessing of γ_i is complete. Thus the overall space requirement of the algorithm is bounded by $O(\log c \cdot (\mathsf{near}(G) + c) + c \log n)$. □

By Proposition 2, we get an ABP P computing a polynomial $p_G(X, Y)$ such that $\mathsf{C\text{-}perm}(G) = \sum_{e_1,\ldots,e_m \in \{0,1\}} p_G(X, e)$, $m = O(c \log n)$. Combining the above algorithm with the closure property of algebraic branching programs over read-once variables given by Proposition 1, we get a non-commutative arithmetic branching program computing $\mathsf{C\text{-}perm}(G)$. It can be seen that size of the resulting branching program is at most $m2^{O((c \log c + \log c \cdot \mathsf{near}(G)) + c \log n)} = n^{O(c)} \cdot c^{\mathsf{near}(G)}$ for large enough n. □

Corollary 3.1. *Let G be as in Theorem 3.1. There is an ABP of size $n^{O(c)}c^{\mathsf{near}(G)}$ computing the Cayley determinant of G.*

Proof. The argument is the same as in Theorem 3.1 except that now the non-deterministic algorithm given in the proof of Theorem 3.1 also needs to compute the sign of the monomial being output. Let $C_1, \ldots, C_r$ be the strongly connected components of G. Then the sign of the permutation corresponding to a cycle cover τ of G is the product of signs of the corresponding cycle covers of C_i. Thus it is enough to modify the algorithm given in the proof of Theorem 3.1 to output the sign of the cycle cover chosen for C_i, the remaining arguments are the same. □

4. Unconditional Lower Bound

We now show that any branching program computing the non-commutative permanent of directed graphs with component size 2 must be of exponential size. This shows that the upper bound in Theorem 3.1 is tight up to a constant factor in the exponent, however, with a different but related parameter. All our lower bound results hold for free algebras over any field $\mathbb{K}$.

Our proof crucially depends on Nisan's ([19]) partial derivative technique. We begin with some notations following his proof. Let f be a non-commutative degree d polynomial in n variables. Let $B(f)$ denote the smallest size of a non-commutative ABP computing f. For $k \in \{0, \ldots, d\}$ let $M_k(f)$ be the matrix with rows indexed by all possible sequences containing k variables and columns indexed by all possible sequences containing $d - k$ variables (repetitions allowed). The entry of $M_k(f)$ at $(x_{i_1} \ldots x_{i_k}, x_{j_1} \ldots x_{j_{d-k}})$ is the coefficient of the monomial $x_{i_1} \cdots x_{i_k} \cdot x_{j_1} \cdots x_{j_{d-k}}$ in f. Nisan established the following result:

Theorem 4.1. *[19] For any homogeneous polynomial f of degree d,*

$$B(f) = \sum_{k=0}^{d} \mathsf{rank}(M_k(f)).$$

We prove lower bounds for the Cayley permanent of graphs with every strongly connected component of size exactly 2, i.e., each strongly connected component being a two-cycle with self loops on the vertices. Note that any collection of $n/2$ vertex disjoint two-cycles can be viewed as a permutation $\pi \in S_n$ consisting of disjoint transpositions and that π is an involution. Conversely, any involution π on n elements represents a graph G_π with connected component size 2.

For a permutation $\pi \in S_n$ let the *cut at i* denoted by $C_i(\pi)$ be the set of pairs $(j, \pi(j))$ that cross i, i.e., $C_i(\pi) = \{(j, \pi(j)) \mid i \in [j, \pi(j)] \cup [\pi(j), j]\}$. The *cut* parameter $\mathsf{cut}(\pi)$ of π is defined as $\mathsf{cut}(\pi) = \max_{1 \le k \le n} |C_k(\pi)|$. Let G be a collection of vertex disjoint 2-cycles denoted by $(a_1, b_1), \ldots, (a_{n/2}, b_{n/2})$ where n is even. The corresponding involution is $\pi_G = (a_1, b_1) \cdots (a_{n/2}, b_{n/2})$. By abusing the notation a bit, we let $\mathsf{cut}(G) = \mathsf{cut}(\pi_G)$. Without loss of generality, assume that $a_i < b_i$, and $a_1 < a_2 < \cdots < a_{n/2}$. Firstly, we note that $\mathsf{cut}(\pi)$ is bounded by $\mathsf{near}(G)$.

Lemma 4.1. *For any collection of disjoint 2-cycles G on n vertices, $\mathsf{cut}(\pi) \le \mathsf{near}(G)$ where π is the involution represented by G.*

Proof. Suppose $\mathsf{cut}(\pi) = r$, and $1 \le i \le n$ be such that $|C_i(\pi)| = r$. Let $(a, b) \in G$ where a is the least value with $a < i$ and $b > i$ be the maximum such value. Then $b - a \ge 2|C_i(\pi)|/2 = r$. This concludes the proof. □

Further, we note that the upper bound given in Theorem 3.1 holds true even if we consider $\mathsf{cut}(G)$ instead of $\mathsf{near}(G)$.

Lemma 4.2. *Let G be a collection of disjoint 2-cycles and self loops where every edge is labeled by a distinct variable or a constant from R. Then there is an ABP of size $2^{O(\mathsf{cut}(G))}n^2$ computing the Cayley permanent of G.*

Proof. The algorithm is the same as in Theorem 3.1. We only need to argue the space bound as in the claim in Theorem 3.1. First note that either $a_i = i$, or i has already occurred in one of the involutions $(a_a, b_1), \ldots, (a_{i-1}, b_{i-1})$. When the algorithm processes the component corresponding to the involution (a_i, b_i), it needs to remember the outgoing edge chosen for b_i (either the self loop or the edge $b_i \to a_i$). Thus at any stage, the number of edges that needs to be stored is bounded by $t = \max_k |C_k(\sigma)|$. The rest of the arguments are exactly the same as in Theorem 3.1. □

Lemma 4.3. *Let G be a collection of ℓ disjoint two-cycles described by the involution π and self loops at every vertex with edge labeled by distinct variables. Then $M_\ell(\mathsf{C\text{-}perm}(G))$ contains $I_2^{\otimes t}$ as a sub-matrix where $t = \max_k |C_k(\pi)|$, $A^{\otimes t}$ is the tensor product of A with itself t times and I_2 is the 2×2 identity matrix.*

Proof. Let $k \in [\ell]$, and $m = |C_k(\pi)| \leq \ell$. Let $C_k(\pi) = \{(a_{i_1}, b_{i_1}), \ldots, (a_{i_m}, b_{i_m})\}$ be such that $a_{i_j} \leq k \leq b_{i_j}$ for all j. Let G_k be the graph restricted to involutions in $C_k(\pi)$. By induction on m, we argue that $M_m(\mathsf{C\text{-}perm}(G_k))$ contains $I_2^{\otimes m}$ as a sub-matrix. The lemma would then follow since $M_m(\mathsf{C\text{-}perm}(G_k))$ is itself a sub-matrix of $M_\ell(\mathsf{C\text{-}perm}(G))$.

We begin with $m = 1$ as the base case. Consider the transposition (a_{i_j}, b_{i_j}), with $a_{i_j} \leq k \leq b_{i_j}$. The corresponding two cycle has four edges. Let f_{i_j} be the Cayley permanent of this graph then $M_1(f_{i_j})$ has the 2×2 identity matrix as a sub-matrix. Let us dwell on this simple part. For ease of notation let the variables corresponding to the self loops be given by x_a, x_b for (a_{i_j}, a_{i_j}) and (b_{i_j}, b_{i_j}) respectively and the edge (a_{i_j}, b_{i_j}) by $x_{(a,b)}$ and the edge (b_{i_j}, a_{i_j}) by $x_{b,a}$. Now our matrix has monomials $x_a, x_{a,b}$ as rows and $x_b, x_{b,a}$ as columns. We can ignore the other orderings as these will always be zero. As the valid cycle covers are given by $x_a x_b$ and $x_{a,b} x_{b,a}$ the proof is clear.

For the induction step, suppose $m > 1$. Suppose $a_1 < a_2 < \cdots < a_m$. Let G'_k be the graph induced by $C_k(\pi) \setminus (a_1, b_1)$. Let $M' = M_{m-1}(\mathsf{C\text{-}perm}(G'_k))$. The rows of M' are labeled by monomials consisting of variables with first index $\leq k$ and the columns of M' are labeled by monomials consisting only of variables with first index $> k$. Let $M = M_m(\mathsf{C\text{-}perm}(G_k))$. M can be obtained from M' as follows: Make two copies of the row labels of M', the first one with monomials pre-multiplied by x_{a_1,a_1}, and the second pre-multiplied by x_{a_1,b_1}. Similarly, make two copies of the columns of M', the first by inserting x_{b_1,b_1} to the column labels of M' at appropriate position, and then inserting x_{b_1,a_1} similarly. Now, the matrix M can be viewed as two copies of M' that are placed along the diagonal. Thus $M = M' \otimes I_2$, combining this with Induction Hypothesis completes the proof. □

Remark 4.1. *It should be noted that the ordering of the variables is crucial in the above argument. If $a_1, b_1 < k$ in the above, then $\mathsf{rank}(M) = \mathsf{rank}(M')$.*

Theorem 4.2. *Let G be a collection of disjoint two cycles described by the involution π and self loops at every vertex, with edges labeled by distinct variables. Then any non-commutative ABP computing the Cayley permanent on G has size at least $2^{\Omega(\mathsf{cut}(G))}$.*

Proof. It is enough to argue that for every k, there is an ℓ with $\mathsf{rank}(M_\ell(f)) \geq 2^{\Omega(|C_k(\pi)|)}$, then the claim follows from Theorem 4.1 as the lower bound is given by the sum over all ℓ. Let $\ell = |C_k(\pi)|$, and suppose $(a_{i_1}, b_{i_1}), \ldots, (a_{i_\ell}, b_{i_\ell})$ are the transpositions crossing k. Let G' be the sub-graph of G induced by the vertices corresponding to the transpositions above. Let $f = \mathsf{C\text{-}perm}(G')$. Applying Lemma 4.3 on G' we conclude that $M_\ell(f)$ has $I_2^{\otimes |C_k(\pi)|}$ as a sub-matrix, i.e., the identity matrix of dimension $2^{|C_k(\pi)|} \times 2^{|C_k(\pi)|}$. Note that f can be obtained by setting weights of the self loops of vertices not in G' to zero, and setting the remaining variables to 1. Moreover, the matrix $M_\ell(f)$ is a sub matrix of $M_\ell(\mathsf{C\text{-}perm}(G))$ obtained by relabeling the rows and columns as per the substitution mentioned above, and removing rows and columns that are zero. We conclude $\mathsf{rank}(M_\ell(\mathsf{C\text{-}perm}(G))) \geq \mathsf{rank}(M_\ell(f)) \geq 2^{|C_k(\pi)|}$. □

Let $\pi = (a_1, b_1) \cdots (a_{n/2}, b_{n/2})$, $a_1 < a_2 < \cdots < a_{n/2}$ be an involution. Then G_π is the set of 2-cycles $(a_1, b_1), \ldots, (a_{n/2}, b_{n/2})$ and self loops at every vertex.

Corollary 4.1. *Let G be a collection of disjoint two cycles described by the involution π and self loops at every vertex, with edges labeled by distinct variables. Then $B(\mathsf{C\text{-}perm}(G)) \in 2^{\Theta(\mathsf{cut}(G))}$. Further, there exists a graph G with $\mathsf{cut}(G) = \Theta(n)$.*

Proof. The first part follows immediately from Lemma 4.2 anf Theorem 4.2. For the second statement, consider the involution π with $\pi(i) = n/2 + i$. It can be seen that $\max_k |C_k(\pi)| = n/2$. Thus by Theorem 4.2 we have any ABP computing the Cayley permanent of G is of size $2^{\Omega(n)}$. Since $\mathsf{cut}(\pi) \leq n$ for any graph, by Lemma 4.2 the result follows. □

Finally, we have,

Theorem 4.3. *For all but a $1/\sqrt{n}$ fraction of graphs G with connected component size 2, any ABP computing the $\mathsf{C\text{-}perm}$ on G requires size $2^{\Omega(n)}$.*

As before, let $n = 2m$ be even. Then an involution π on $\{1, \ldots, n\}$ with $\pi(i) \neq i$ represents a collection of m intervals

$$I_\pi = \{[i, \pi(i)] \mid 1 \leq i \leq n,\ i < \pi(i)\}.$$

Let H_π be the interval graph formed by the intervals in I_π.

Lemma 4.4. *Let π be an involution and H_π be the interval graph as defined above. Then $\mathsf{cut}(\pi) \geq l/n$ where l is the number of edges in H_π.*

Proof. For every edge (a, b) in H_π, the corresponding intervals $I_a = [i, \pi(i)]$ and $I_b = [j, \pi(j)]$ have non empty intersection. Suppose $i < j$, then $j \in [i, \pi(i)]$. (In the case when $j > \pi(j)$, we have $\pi(j) \in [i, \pi(i)]$. Other cases can be handled analogously.) Thus every edge in H_π contributes at least one distinct interval $[i, \pi(i)]$ with $i \leq k \leq \pi(i)$, i.e., it contributes a value to $C_k(\pi)$. Then $l \leq \sum_k |C_k(\pi)|$. This concludes the proof. □

Scheinerman [21] showed that, random interval graphs have $\Omega(n^2)$ edges with high probability, i.e.,

Theorem 4.4. *[21] Let H_π be an interval graph where π is an involution on $[n]$ chosen uniformly at random. Then H_π has at least $n^2/3 - n^{7/4}$ edges with probability at least $1 - 1/\sqrt{n}$.*

Corollary 4.2. *For an involution π on $[n]$ chosen uniformly at random, we have $\mathsf{cut}(\pi) = \Omega(n)$ with probability $1 - 1/\sqrt{n}$.*

Theorem 4.3 now follows.

5. #P completeness

In this section, we show multiple hardness results for simple polynomials over certain classes of non-commutative algebras. We give a #P completeness result for specific graphs of component size at most six. The completeness result is obtained by a careful

analysis of the parameters in the reduction from $\#SAT$ to non-commutative determinant given recently by Gentry [14] and the small modification we will do to make this proof work for the Cayley Permanent.

Theorem 5.1. *Let R be a division algebra over a field $\mathbb{K}$ of characteristic zero containing the algebra of 2×2 matrices over $\mathbb{K}$. Computing the Cayley Permanent on graphs with component size 6 with edges labeled from R is $\#\mathsf{P}$ complete.*

Proof. It is known that counting the number of satisfying assignments in a 2-CNF formula where every variable occurs at most three times is already $\#\mathsf{P}$ complete ([20]). Let ϕ be a 2-CNF where every variable occurs at most three times with k clauses. We complete the proof by a careful analysis of the reduction given in Theorem 6 of [14] applied to ϕ.

Lemma C.1 gives a product program of length $2^2+2^{2-1}-2 = 4$ for computing a disjunction of two literals. In fact the program for $x_1 \vee x_2$ is given by $(1,(s,I_2)),(2,(r,I_2)),(1,(s,I_2)),(2,(r^{-1},I_2))$ where I_2 is the 2×2 identity matrix in R.

Let t be a 2×2 matrix as in Theorem C.1, namely t has one in the upper left corner and zeros elsewhere. Suppose P_c is the product program as given above for the clause indexed by c for $1 \leq c \leq k$. Then product program for ϕ is given by

$$P = \left(\prod_{1 \leq c \leq k} t \cdot P_c \right) t.$$

This immediately shows if every variable occurs at most three times in ϕ, the product program above reads a bit of the input at most 6 times. Let $\mathcal{I}_\ell$ as before where $1 \leq \ell \leq 4k$. We have $\mathcal{I}_\ell \leq 6$ for all $1 \leq \ell \leq 4k$ by the above argument. Let $\mathcal{I}_\ell$ have the elements $i_{\ell,1}, \ldots, i_{\ell,|\mathcal{I}_\ell|}$. Let π_0 be the identity permutation and $\pi_1(i_{\ell,\kappa}) = i_{\ell,\kappa+1 \mod |\mathcal{I}_\ell|}$. Define the following permuted "block barber pole" ([14]) matrix.

$$M[i,j] = \begin{cases} a_{i,b} & \text{if } j = \pi_b(i) \\ 0 & \text{otherwise.} \end{cases}$$

While Gentry has in the cell $a_{i_{l,1},i_{l,2}}$ a different factor, $(-1)^{|\mathcal{I}_\ell|-1}$ to be precise, this factor is only canceling the sign of the determinant as can be seen in the proof of Theorem C.2. As the rows and columns for $\mathcal{I}_i, \mathcal{I}_j$ for $i \neq j$ are disjunct this matrix corresponds to cycles of length $|\mathcal{I}_\ell| \leq 6$. This concludes the proof. □

It is known that computing the commutative permanent of the weighted adjacency matrix of a planar graph is as hard as the general case [10]. We observe that the reduction in [10] extends to the non-commutative case.

Theorem 5.2. $\mathsf{C\text{-}perm} \leq_m^p \mathsf{planar{-}C\text{-}perm}$*; and* $\mathsf{C\text{-}det} \leq_m^p \mathsf{planar{-}C\text{-}det}$*. Moreover, the above reductions work over any non-commutative algebra.*

Proof. The proof is essentially the same as in [10]. We give a brief sketch here for the sake of completeness. Let G be a weighted digraph. Consider an arbitrary embedding $\mathcal{E}$ of G. Obtain a new graph by changing the graph as follows:

- For each pair of edges (u, v) and (u', v') that cross each other in the embedding $\mathcal{E}$, do the following:
- introduce two new vertices a and b; and
- new edges $\{(a, b), (b, a), (u', a), (a, v), (u, b), (b, v')\}$ replacing (u, v) and (u', v').

Note that any of the iterations above do not introduce any new crossings, and hence the process terminates after at most $O(n^2)$ many steps where n is the number of vertices in G. Weight of (u, v) is given to (v, a) and (u', v') is given to (v', b). The remaining edges get the weight 1. By the construction, we can conclude that $\mathsf{C\text{-}perm}(G) = \mathsf{C\text{-}perm}(G')$ and $\mathsf{C\text{-}det}(G) = \mathsf{C\text{-}det}(G')$. □

We demonstrate some more families of polynomials whose commutative variants are easy but certain non-commutative variants are as hard as the permanent polynomial. We begin with a non-commutative variant of the elementary symmetric polynomial. The elementary symmetric polynomial of degree d, $\mathsf{Sym}_{n,d}$ is given by $\mathsf{Sym}_{n,d}(x_1, \ldots, x_n) = \sum_{S \subseteq [n],\ |S|=d} \prod_{i \in S} x_i$. There are several non-commutative variants of the above polynomial. The first one is analogous to the Cayley permanent, i.e., $\mathsf{Cayley{-}Sym}_{n,d} = \sum_{S=\{i_1 < i_2 < \cdots < i_d\}} \prod_{j=1}^{d} x_{i_j}$. It is not hard to see that the above mentioned non-commutative version of $\mathsf{Cayley{-}Sym}_{n,d}$ can be computed by depth 3 non-commutative circuits for every value of $d \in [n]$. However, the above definition is not satisfactory, since it is not invariant under permutation of variables, which is the inherent property of elementary symmetric polynomials. We define a variant of non-commutative elementary symmetric polynomial which is invariant under the permutation of variables.

$$\mathsf{nc{-}Sym}_{n,d}(x_1, \ldots, x_n) \stackrel{\triangle}{=} \sum_{\{i_1, \ldots, i_d\} \subseteq [n]} \sum_{\sigma \in S_d} \prod_{j=1}^{d} x_{i_{\sigma(j)}}.$$

We show that with coefficients from the algebra of $n \times n$ matrices allowed, $\mathsf{nc{-}Sym}_{n,d}$ cannot be computed by polynomial size circuits unless $\mathsf{VP} = \mathsf{VNP}$. We need the following definition introduced in [2, 4].

Definition 5.1. *The* Hadamard product *between two polynomials $f = \sum_m \alpha_m m$ and $g = \sum_m \beta_m m$, written as $f \odot g$, is defined as $f \odot g = \sum_m \alpha_m \beta_m m$.*

Theorem 5.3. *Over any $\mathbb{K}$ algebra R containing the $n \times n$ matrices as a sub-algebra, $\mathsf{nc{-}Sym}_{n,n}$ does not have polynomial size arithmetic circuits unless $\mathsf{perm}_n \in \mathsf{VP}$.*

Proof. Suppose that $\mathsf{nc{-}Sym}_{n,n}$ has a circuit C of size polynomial in n. We need to show that $\mathsf{perm} \in \mathsf{VP}$. Let $X = (x_{i,j})_{1 \le i,j \le n}$ be matrix of variables, and $y_1, \ldots, y_n$ be distinct variables different from $x_{i,j}$. In the commutative setting, it was observed in [24] that $\mathsf{perm}(X)$ equals the coefficient of $y_1 \cdots y_n$ in the polynomial

$$P(X, Y) \stackrel{\triangle}{=} \prod_{i=1}^{n} \left(\sum_{j=1}^{n} x_{i,j} y_j \right) \tag{1}$$

over the polynomial ring $\mathbb{K}[x_{1,1}, \ldots, x_{n,n}]$. However, the same cannot be said in the case of non-commuting variables. If $x_{i,j}y_k = y_k x_{i,j}$ for $i, j, k \in [n]$, then in the non-commutative development of (1), the sum of coefficients of *all* permutations of the monomial $y_1 \cdots y_n$ equals $\mathsf{perm}(X)$ i.e., the commutative permanent. Hence the value $\mathsf{perm}(X)$ can be extracted using a Hadamard product with $\mathsf{nc\text{-}Sym}_{n,n}(y_1, \ldots, y_n)$ and then substituting $y_1 = 1, \ldots, y_n = 1$. However, we cannot assume $x_{i,j}y_k = y_k x_{i,j}$, since the Hadamard product may not be computable under this assumption. Let $\ell = \sum_{i,j} x_{i,j}$. Now we argue that $\mathsf{perm}(X) = (\mathsf{nc\text{-}Sym}_{n,n}(\ell y_1, \ldots, \ell y_n) \odot P)(y_1 = 1, \ldots, y_n = 1)$. Given a permutation $\sigma \in S_n$, there is a unique monomial $m_\sigma = x_{1,\sigma(1)}y_{\sigma(1)} \cdots x_{n,\sigma(n)}y_{\sigma(n)}$ in P containing the variables $y_{\sigma(1)}, \ldots, y_{\sigma(n)}$ in that order. Thus taking Hadamard product with P filters out all monomials but m_σ from the term $\prod_{i=1}^{n} \ell y_{\sigma(i)}$. The monomials where a y_j occurs more than once are eliminated by $\mathsf{nc\text{-}Sym}_{n,n}(\ell y_1, \ldots, \ell y_n)$. Thus the only monomials that survive in the Hadamard product are of the form m_σ, $\sigma \in S_n$. Now substituting $y_i = 1$ for $i \in [n]$ we get $\mathsf{perm}(X) = (\mathsf{nc\text{-}Sym}_{n,n}(\ell y_1, \ldots, \ell y_n) \odot P)(y_1 = 1, \ldots, y_n = 1)$.

Note that the polynomial $P(X, Y)$ can be computed by an ABP of size $O(n^2)$. Then, by [2, 4], we obtain an arithmetic circuit D of size $O(n^2\mathsf{size}(C))$ that computes the polynomial $\mathsf{nc{-}Sym}_{n,n} \odot P$. Substituting $y_1 = 1, \ldots, y_n = 1$ in D gives the required arithmetic circuit for $\mathsf{perm}(X)$. □

Further, let $\mathsf{snc\text{-}Sym}_{n,n}(x_1, \ldots, x_n) \stackrel{\triangle}{=} \sum_{\sigma \in S_n} \mathsf{sgn}(\sigma) \prod_{i=1}^{n} x_{i,\sigma(i)}$. We have

Corollary 5.1. *Over a $\mathbb{K}$ algebra containing the algebra of $n \times n$ matrices,* $\mathsf{snc\text{-}Sym}_{n,n}$ *is not in* VP *unless* C-det *has polynomial size arithmetic circuits.*

While similar, this result is unrelated with the result of [15] as their ordered polynomials is more related to $\mathsf{Cayley{-}Sym}_{n,d}$ than $\mathsf{nc{-}Sym}_{n,n}$.

Barvinok [6] showed that computing the permanent of an integer matrix of constant rank can be done in strong polynomial time. In a similar spirit, we explore the complexity of computing the Cayley permanent of bounded rank matrices with entries from $\mathbb{K} \cup \{x_1, \ldots, x_n\}$. We consider the following notion of rank for matrices with variable entries. Let $A \in (\mathbb{K} \cup \{x_1, \ldots, x_n\})^{n \times n}$. Then $\mathsf{row\text{-}rank}(A) = \max_{a_1, \ldots, a_n \in \mathbb{K}} \mathsf{rank}(A|_{x_1 = a_1, \ldots, x_n = a_n})$. The column rank of A is defined analogously. As opposed to the case of the commutative permanent, for any algebra R containing the algebra of $n \times n$ matrices over $\mathbb{K}$, we have:

Corollary 5.2. C-perm *and* C-det *of rank one matrices with entries from $\mathbb{K} \cup \{x_1, \ldots, x_n\}$ over any $\mathbb{K}$ algebra does not have polynomial size arithmetic circuits unless* $\mathsf{perm} \in \mathsf{VP}$.

Proof. We will argue the case of C-perm. Let $x_1, \ldots, x_n$ be non-commuting variables. Consider the matrix A with $A[i, j] = x_j$, $1 \leq i, j \leq n$. A has rank one over $\mathbb{K}$. We then have $\mathsf{nc{-}Sym}_{n,n}(x_1, \ldots, x_n) = \mathsf{C\text{-}perm}(A)$. The result now follows by applying Theorem 5.3. For C-det, we can use Corollary 5.1 in place of Theorem 5.3 in the argument above. □

6. Computational problems on non-commutative circuits

Computing Coefficients

In this section we consider various computational problems on arithmetic circuits, restricted to the non-commutative setting. We start with the problem of computing the coefficient of a given monomial in the polynomial computed by an arithmetic circuit. In the commutative setting, the problem lies in the second level of the counting hierarchy [13] and is known to be hard for $\#P$ [18]. It was first seen in [3] that mcoeff is easy to compute in the non-commutative case. We provide a different proof of the fact as it is useful in the arguments used later in this section.

Problem 1 (Monomial Coefficient(mcoeff))**.** Input*: A non-commutative arithmetic circuit C, a non-commutative monomial m of degree d.*
Output*: The coefficient of monomial m in the polynomial computed by C.*

Theorem 6.1. *[3]* mcoeff *is in* P*.*

Proof. Suppose that the monomial $m = x_{j_1} \cdots x_{j_d}$ and is given as an ordered listing of variables. Let f be a non-commutative polynomial. Then we have the following recursive formulation for the coefficient function $\mathsf{mc} : \mathbb{K}\{x_1, \ldots, x_n\} \times \mathcal{M} \to \mathbb{K}$, where $\mathcal{M}$ is the set of all non-commutative monomials in variables $\{x_1 \ldots, x_n\}$.

$$\mathsf{mc}(f, m) = \begin{cases} \alpha & \text{if } f = \alpha y_j \text{ and } m = y_j, \\ 0 & \text{if } f = \alpha y_j \text{ and } m = y_i, i \neq j \ , \\ \mathsf{mc}(g, m) + \mathsf{mc}(h) & \text{if } f = g + h, \\ \sum_{\ell=1}^{d+1} \mathsf{mc}(g, m_\ell) \times \mathsf{mc}(h, m'_\ell) & \text{if } f = g \times h. \end{cases} \tag{2}$$

where $m_\ell = x_{i_1} \cdots x_{i_{\ell-1}}$ and $m'_\ell = x_{i_\ell} \cdots x_{i_d}$. However, if we apply the above recursive definition on the circuit C in a straightforward fashion, the time required to compute $\mathsf{mc}(f, m)$ will be $d^{O(\mathsf{depth}(C))}$, since $\mathsf{depth}(C)$ could be as big as $\mathsf{size}(C)$, the running time would be exponential. However, we can have a more careful implementation of the above formulation by allowing a little more space.

For $\ell < k \in [1, d]$, let $m_{\ell,k} = x_{i_{\ell+1}} \cdots x_{i_k}$, and $M = \{m_{\ell,k} \mid 0 \leq \ell \leq d-1, 0 \leq k \leq d\}$. Consider a gate v in the circuit C. Note that in the process of computing $\mathsf{mc}(f, m)$, we require only the values from the set $M_v = \{\mathsf{mc}(p_v, m') \mid m' \in M\}$, where p_v is the polynomial computed at v. Thus it is enough to compute and maintain the values $\mathsf{mc}(p_v, m'), m' \in M$ in a bottom up fashion. For the base case, compute the values for polynomials computed at a leaf gate v as follows, let $m' \in M$ and $\alpha \in R$

$$\mathsf{mc}(p_v, m') = \begin{cases} \alpha & \text{if } p_v = \alpha \in R, \text{ and } m' = \emptyset, \\ \alpha & \text{if } p_v = \alpha x_j, \alpha \in R, \text{ and } m' = x_j, \\ 0 & \text{otherwise.} \end{cases}$$

For other nodes, we can apply the recursive formula given in (2). If $p_v = p_{v_1} + p_{v_2}$, then the value $\mathsf{mc}(p_v, m')$ can be computed using (2) as the values $\mathsf{mc}(p_{v_1}, m')$ and $\mathsf{mc}(p_{v_2}, m')$

are available by induction. If $p_v = p_{v_1} \times p_{v_2}$, then by induction, the values $\mathsf{mc}(p_{v_i}, m'')$ are available for prefix and suffix of the monomial m', as every such monomial occurs as $m_{i,j \in M}$ for some $i < j$. Now, $\mathsf{mc}(p_v, m')$ can be computed by (2). For the space bound, the algorithm uses $O(d^2)$ registers for each gate in C and hence the overall space used is $O(d^2\mathsf{size}(C))$ many registers. For a given monomial $m' \in M$, at most d arithmetic operations are required in the worst case. Thus, the number of arithmetic operations is bounded by $O(d^3\mathsf{size}(C))$. □

Coefficient function as a polynomial

In the commutative setting, the coefficient function of a given polynomial can be represented as a polynomial [18]. Thus it is desirable to study the arithmetic circuit complexity of coefficient functions. However, over non-commutative rings, we need a carefully chosen representation of monomials to obtain an arithmetic circuit that computes the coefficient function for a given polynomial with small circuits. In the proof of Theorem 6.1, we have used an ordered listing of variables as a representation of the monomial m. Here we use a vector representation for non-commutative monomials of a given degree d. Let $Y = \{y_{1,1}, \ldots, y_{1,n}, y_{2,1}, \ldots y_{d,n}\}$ be a set of nd distinct variables, and let $\tilde{Y}_i = (y_{i,1}, \ldots, y_{i,n})$. The vector of variables $\tilde{Y}_\ell$ can be seen as representing the characteristic vector of x_{i_ℓ}, i.e., $y_{i,i_\ell} = 1$, and $y_{i,j} = 0, \forall j \neq i_\ell$. In essence, $y_{i,j}$ stands for the variable x_j at the i-th position in the monomial. Let $f(x_1, \ldots, x_n)$ be a polynomial of degree d, then we can define the coefficient polynomial $\mathsf{pc}_f(Y)$ as

$$\mathsf{pc}_f(Y) = \sum_{D=1}^{d} \sum_{(i_1,\ldots,i_D)\in[n]^D} \prod_{\ell=1}^{D} \left[\mathsf{mc}(f, x_{i_1} \cdots x_{i_D}) y_{\ell,i_\ell} \prod_{j\neq k} (1 - y_{\ell,j} y_{\ell,k}) \right].$$

Theorem 6.2. *For any non-commutative polynomial f that can be computed by a polynomial size arithmetic circuit, $\mathsf{pc}_f(Y)$ has a polynomial size arithmetic circuit.*

Proof of Theorem 6.2. We will apply (2) to obtain an arithmetic circuit computing the polynomial $\mathsf{pc}_f(Y)$. Let C be an arithmetic circuit of size s, computing f. By induction on the structure of C, we construct a circuit C' for $\mathsf{pc}_f(Y)$. Note that, it is enough to compute homogeneous degree D components $[\mathsf{pc}_f(Y)]_D$ of $\mathsf{pc}_f(Y)$, where

$$[pc_f(Y)]_D = \sum_{(i_1,\ldots,i_D)\in[n]^D} \prod_{\ell=1}^{D} \left[\mathsf{mc}(f, x_{i_1} \cdots x_{i_D}) y_{\ell,i_\ell} \prod_{j\neq k} (1 - y_{\ell,j} y_{\ell,k}) \right].$$

Let $Y^{i,j}$ denote the set of variables in the vectors $\tilde{Y}_{i+1}, \ldots, \tilde{Y}_j$. In the base case, we have $C = \gamma \in \{x_1, \ldots, x_n\} \cup R$. Then the all of the homogeneous components of $\mathsf{pc}_f(Y)$ can

be described as follows.

$$[\mathsf{pc}_f(Y)]_0 = \begin{cases} \gamma & \text{if } Y = \emptyset \text{ and } \gamma \in R \\ 0 & \text{otherwise.} \end{cases}$$

$$[\mathsf{pc}_f(Y)]_1 = \begin{cases} 1 & \text{if } Y = e_j \text{ and } \gamma = x_j \\ 0 & \text{if } Y = \emptyset \text{ and } \gamma \in R \\ 0 & \text{otherwise.} \end{cases}$$

$$[\mathsf{pc}_f(Y)]_{i>1} = 0.$$

Naturally, the induction step has two cases: $f = g + h$ and $f = g \cdot h$.
Case 1: $f = g + h$, then for any D

$$[\mathsf{pc}_f(Y)]_D = [\mathsf{pc}_g(Y)]_D + [\mathsf{pc}_h(Y)]_D \quad \forall\ D.$$

Case 2: $f = g \times h$, then for any D

$$[\mathsf{pc}_f(Y)]_D = \sum_{i=0}^{d} \sum_{j=0}^{D} [\mathsf{pc}_g(Y^{1,i})]_j [\mathsf{pc}_h(Y^{i+1,d})]_{D-j}$$

where $Y = \tilde{y}_1, \ldots, \tilde{y}_d$. The size of the resulting circuit C' is $O(d^3\mathsf{size}(C))$, and C' can in fact be computed in time $O(d^3\mathsf{size}(C))$ given C as the input. □

Partial Coefficient functions

For a given commutative polynomial let $f(X) = \sum_m c_m m$, the partial coefficient of a given monomial m ([18]) is a polynomial defined as $\mathsf{pcoeff}(f, m) = \sum_{m', m|m'} c_{m'} \frac{m'}{m}$.

We extend the above definition to the case of non-commutative polynomials as follows. Let f be non-commutative polynomial, and m a non-commutative monomial. Then $\mathsf{pcoeff}(f, m) = \sum_{m'=m\cdot m''} c_{m'} m''$.

The corresponding computational problem can be defined in the following way.

Problem 2 (Coefficient Polynomial (pcoeff))**.** Input*: A non-commutative arithmetic circuit C computing a polynomial f, and a monomial m.*
Output*: A non-commutative arithmetic circuit that computes* $\mathsf{pcoeff}(f, m)$*.*

Theorem 6.3. pcoeff *can be computed in deterministic time* $\mathsf{poly}(\mathsf{size}(C), n, \deg(m))$.

Proof. The algorithm is similar to the proof of Theorem 6.1, except that we need to construct an arithmetic circuit rather than a value. We use the following recursive formulation similar to (2).

If $f = \alpha \in R \cup \{x_1, \ldots, x_n\}$ and $m = \emptyset$ then $\mathsf{pcoeff}(f, m) = \alpha$. For the summation $f = g + h$ we compute $\mathsf{pcoeff}(f, m) = \mathsf{pcoeff}(g, m) + \mathsf{pcoeff}(h, m)$. The final case to handle is a multiplication gate. We define shorthand for sets of variables. Let $m = x_1 \cdots x_d$, $m_i = x_1 \cdots x_i$ and $m'_i = x_{i+1} \cdots x_d$ the rest of the monomial. We define $m_0 = \emptyset$. Then

$\mathsf{pcoeff}(f,m) = \sum_{i=0}^{d-1} \mathsf{mc}(g,m_i)\mathsf{pcoeff}(h,m_i') + \mathsf{pcoeff}(g,m) \cdot \mathsf{pcoeff}(h,\emptyset)$. The rest of the proof is analogous to that of Theorem 6.1 except that, we need to compute and store the values $\mathsf{mc}(p_v, m_{i,j})$, and $\mathsf{pcoeff}(p_v, m_{i,j})$ for every gate v in the circuit in a bottom up fashion. □

Acknowledgements: The authors like to thank V. Arvind and Markus Bläser for helpful discussions and pointing out specific problems to work on. The authors also thank anonymous referees for their comments which helped in improving the presentation. This work was partially done while the first author was visiting IIT Madras sponsored by the Indo-Max-Planck Center for Computer Science.

References

[1] S. Arora and B. Barak. *Computational Complexity: A Modern approach*. Cambridge Univeristy Press, 2009.

[2] V. Arvind, P. S. Joglekar, and S. Srinivasan. Arithmetic circuits and the hadamard product of polynomials. In *FSTTCS*, pages 25–36, 2009.

[3] V. Arvind, P. Mukhopadhyay, and S. Srinivasan. New results on noncommutative and commutative polynomial identity testing. CCC, pages 268–279, 2008.

[4] V. Arvind and S. Srinivasan. On the hardness of the noncommutative determinant. In *STOC*, pages 677–686, 2010.

[5] H. Aslaksen. Quaternionic determinants. *The Math. Int.*, 18(3):57–65, 1996.

[6] A. I. Barvinok. Two algorithmic results for the traveling salesman problem. *Mathematics of Operations Research*, 21(1):65–84, 1996.

[7] M. Bläser. Noncommutativity makes determinants hard. In *ICALP (1)*, pages 172–183, 2013.

[8] P. Bürgisser. *Completeness and reduction in algebraic complexity theory*, volume 7. Springer, 2000.

[9] S. Chien and A. Sinclair. Algebras with polynomial identities and computing the determinant. *SIAM Journal on Computing*, 37(1):252–266, 2007.

[10] S. Datta, R. Kulkarni, N. Limaye, and M. Mahajan. Planarity, determinants, permanents, and (unique) matchings. *ToCT*, 1(3):10, 2010.

[11] U. Flarup, P. Koiran, and L. Lyaudet. On the expressive power of planar perfect matching and permanents of bounded treewidth matrices. In *ISAAC*, pages 124–136, 2007.

[12] U. Flarup and L. Lyaudet. On the expressive power of permanents and perfect matchings of matrices of bounded pathwidth/cliquewidth. *ToCS*, 46(4):761–791, 2010.

[13] H. Fournier, G. Malod, and S. Mengel. Monomials in arithmetic circuits: Complete problems in the counting hierarchy. In *STACS*, pages 362–373, 2012.

[14] C. Gentry. Noncommutative determinant is hard: A simple proof using an extension of barrington's theorem. In *CCC*, pages 181–187, June 2014.

[15] P. Hrubes, A. Wigderson, and A. Yehudayoff. Non-commutative circuits and the sum-of-squares problem. In *STOC*, pages 667–676. ACM, 2010.

[16] N. Limaye, G. Malod, and S. Srinivasan. Lower bounds for non-commutative skew circuits. *Electronic Colloquium on Computational Complexity (ECCC)*, 22:22, 2015.

[17] M. Mahajan and B. V. R. Rao. Small space analogues of valiant's classes and the limitations of skew formulas. *Computational Complexity*, 22(1):1–38, 2013.

[18] G. Malod. The complexity of polynomials and their coefficient functions. In *IEEE Conference on Computational Complexity*, pages 193–204, 2007.

[19] N. Nisan. Lower bounds for non-commutative computation (extended abstract). In *STOC*, pages 410–418, 1991.

[20] D. Roth. On the hardness of approximate reasoning. *Artif. Intell.*, 82(1-2):273–302, 1996.

[21] E. R. Scheinerman. Random interval graphs. *Combinatorica*, 8(4):357–371, 1988.

[22] A. Shpilka and A. Yehudayoff. Arithmetic circuits: A survey of recent results and open questions. *FTTS*, 5(3-4):207–388, 2010.

[23] L. G. Valiant. Completeness classes in algebra. STOC '79, pages 249–261, 1979.

[24] J. von zur Gathen. Feasible arithmetic computations: Valiant's hypothesis. *J. Symb. Comput.*, 4(2):137–172, 1987.

A. Permanent as a sum over cycle covers

Let G be a weighted directed graph on n vertices and $x_{i,j}$ denote the weight of the directed edge $(i,j) \in E(G)$, the Cayley permanent of G denoted by $\mathsf{C\text{-}perm}(G)$ is the permanent of the weighted adjacency matrix of G. A cycle cover of G is a collection of vertex disjoint cycles $\pi = (C_1, \ldots C_k)$ that cover all the vertices in G. Cayley weight of a cycle cover is the product of weights of edges in the cover, multiplied in the Cayley order, i.e., if $\pi(i)$ denotes the successor of node i in the cycle cover π, then the weight of the cover π is $\prod_{i=1}^{n} x_{i,\pi(i)}$. It is known that [8] $\mathsf{C\text{-}perm}(G)$ is the sum of the Cayley weights of all cycle covers of G.

B. Proof Sketch for Proposition 1

Note that the proof given in [17] uses the equivalence of skew circuits with ABPs, which does not hold in the non-commutative setting. Our argument is similar to the one in [17] except that we argue over ABPs themselves rather than skew circuits. We give a sketch of the proof here. Let P be an ABP computing the non-commutative polynomial given by $p_P(x_1, \ldots, x_n, y_1, \ldots, y_m)$. Let $0 = i_0 < i_1 < \cdots < i_m$ be the layers of P that witness the fact that P is read-once certified in $Y = \{y_1, \ldots, y_m\}$. Without loss of generality assume that every layer of P has exactly r nodes. Let $g_1^j, \ldots, g_r^j$ be the nodes in the layer i_j. Note that variable y_j is read in layers between i_{j-1} to i_j and is never used beyond that point and no other variable from Y appears in layers between i_{j-1} and i_j. Let P_j be the portion of P consisting only of layers of P from i_{j-1} to i_j. Let $p_j[\ell, k]$ be the polynomial represented as sum of weights of $g_\ell^j \rightsquigarrow g_k^{j+1}$ paths in P. Then

$$p_P = \sum_{k_1,\ldots,k_m} p_0[s,k_1]p_1[k_1,k_2]\cdots p_{m-2}[k_{m-2},k_{m-1}]\cdot p_{m-1}[k_{m-1},t].$$

For $0 \le j < m$, $k, k' \in \{1, \ldots r\}$ and $b \in \{0,1\}$, let $p_j^b[k,k'] = p_j[k,k'](y_{j+1} = b)$ where we substitute y_{j+1} with b. Then

$$\sum_{e_1,\ldots,e_m \in \{0,1\}} p_P(X, e_1, \ldots, e_m)$$

is equal to

$$\sum_{k_1,\ldots,k_m} \big[\, (p_0^0[s,k_1] + p_0^1[s,k_1]) \cdot (p_1^0[k_1,k_2] + p_1^1[k_1,k_2]) \cdot \cdots \\ \cdot (p_{m-2}^0[k_{m-2},k_{m-1}] + p_{m-2}^1[k_{m-2},k_{m-1}]) \\ \cdot (p_{m-1}^0[k_{m-1},t] + p_{m-1}^1[k_{m-1},t]) \,\big].$$

Thus we can take sums of P_j with $y_j = 0$ and $y_j = 1$ independent of $P_{j'}$ when $j \neq j'$. In the following we describe this construction and omit the proof of correctness and the bound on the size of the resulting ABP. Now create two copies P_j^0 and P_j^1 where P_j^b is

obtained by setting $y_j = b$ for $b \in \{0,1\}$. For every $1 \leq k \leq t$, merge the copies of g_k^j in P_j^0 and P_j^1 into a single vertex h_k^j, and similarly copies of g_k^{j+1} in P_j^0 and P_j^1 into a single vertex h_k^{j+1}. Let Q_j be the resulting program. Let Q be the ABP $Q_1 \cdot Q_2 \cdots Q_m$ obtained by doing the following for every j: For $1 \leq k \leq t$ glue the copies of h_k^j in the top layer of Q_j and first layer of Q_{j+1} to get a single vertex. Let Q be the resulting ABP. From the observation above, we have

$$p_Q = \sum_{e_1,\ldots,e_m \in \{0,1\}} p_P(x_1, \ldots, x_n, e_1, \ldots, e_m)$$

as required. By the construction above, size of Q is at most twice that of P. This completes the proof.

C. Recap of Gentry's Proof

For clarity we will repeat the proof of Gentry with some corrections by Goldreich. Readers familiar with Gentry's Proof can skip this section. We will denote the inverse of an element g in a group by g^{-1}.

Definition C.1. *Let R be some algebra. A product program P over R with n instructions for an input of length m is given by $P = (a_0, (\iota_1, a_{1,0}, a_{1,1}), \ldots, (\iota_n, a_{n,0}, a_{n,1}))$. Where we call the sequence of length n of the form $(\iota_i, a_{i,0}, a_{i,1})_{i\in[n]} \in ([m], R, R)^n$ the instructions and a_0 the starting element.*

It computes on an input $x_1, \ldots, x_m \in \{0,1\}$ the product

$$a_0 \cdot \prod_{i=1}^{n} a_{i,x_{\iota_i}}.$$

In words, our product program decides for every instruction $(\iota_i, a_{i,0}, a_{i,1})$ if it should multiply $a_{i,0}$, if the bit of x at the position ι_i is zero, or $a_{i,1}$, if the bit of x at the position ι_i is one. We will generally not distinguish between an input as vector of length m or a string of length m and will index the string x with x_i to mean the ith bit.

Lemma C.1 ([14, Lemma 2])**.** *For any division algebra R, the group of units of $R^{2\times 2}$ contains a subgroup isomorphic to S_3. In particular, $R^{2\times 2}$ contains the matrices*

$$r = \begin{pmatrix} 0 & -1 \\ 1 & -1 \end{pmatrix} \text{ and } s = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

Proof. It is clear that $R^{2\times 2}$ contains these matrices. It is easy to see that r has order three and s order two and hence by Lagrange's Theorem they generate a group of order at least six. We know that the only elements in the group are $1, s, r, r^2, rs, sr$ as $rs = sr^2$ and $r^2 s = sr$ and s has order two. As S_3 is the only non-abelian group of order six the lemma holds. □

With this we can now show a product program that outputs one if an assignment satisfies a 3-CNF formula and zero else.

Lemma C.2 ([14, Lemma 3])**.** *There exists a product program of length $2^d + 2^{d-1} - 2$ over the group S_3 that computes a disjunction of d literals. It outputs 1 if x satisfies the disjunction and r otherwise.*

Proof. Let 1 be the multiplicative neutral element of the $R^{2\times 2}$.

We give a proof by induction and assume a_0 to be 1. Let $d = 1$. As $a_{1,0} = r$ and $a_{1,1} = 1$ the proof is clear. Let us now assume that the lemma is true for $d-1$ and let P_{d-1} be the constructed product program. Let $x_1, \ldots, x_n$ be the bits of our input and $b_0 = s$ and $b_1 = 1$. Then we construct the program such that the multiplication will be performed as follows:

$$b_{x_d} \cdot P_{d-1}(x_1, \ldots, x_{d-1}) \cdot b_{x_d} \cdot (P_{d-1}(x_1, \ldots, x_{d-1}))^{-1}.$$

Here $(P_{d-1}(x_1, \ldots, x_{d-1})^{-1}$ is replacing all instructions $(\iota_i, (\alpha, \beta))$ by the corresponding instruction $(\iota_i, (\alpha^{-1}, \beta^{-1}))$.

Let us now prove the correctness. If $x_i = 1$ then it is clear that the layer i evaluates to one as one commutes with all elements, especially P_{d-1} and P_{d-1}^{-1}. Hence all layers above will also evaluate to one as $b_0^2 = b_1^2 = 1$.

If all bits of the input are zero then $P_{d-1}(x_1, \ldots, x_{d-1}) = r$ by induction, $b_{x_d} = s$ and hence $P_d(x_1, \ldots, x_d) = srsr^{-1}$. By the equalities above this is equal to

$$s(rs)r^{-1} = s(sr^2)r^{-1} = s^2 r = 1r.$$

□

Let

$$t = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}.$$

Theorem C.1 ([14, Theorem 6])**.** *For any division algebra R and any constant d one can construct a product program of length $k(2^d + 2^{d-1} - 2)$ for a d-CNF formula with k clauses. It outputs t if the formula is satisfied by x and 0 otherwise.*

Proof. Let our clauses be given by $c_1, \ldots, c_k$. By Lemma C.2 we get a product program P_{c_i} for every clause c_i. Then we construct our product program P to compute the multiplication as follows:

$$\left(\prod_{i=1}^{k} t \cdot P_{c_i}(x_1, \ldots, x_m) \right) t.$$

Here the multiplication with t can easily be simulated with an instruction of the form $(1, (t, t))$.

For ease of notation we have written the complete variable set for our clause product programs but we can easily remove unneeded variables from a clause.

Let us give a correctness argument. Suppose one of the clauses is not satisfied. This then contributes a value of $t \cdot r$ to our product. It can be seen that this program has only two possible outcomes if one equation is not fulfilled. Namely,

$$t \cdot r = \begin{pmatrix} 0 & -1 \\ 0 & 0 \end{pmatrix}$$

and for any g

$$t \cdot r \cdot t \cdot g = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix},$$

as $t \cdot r \cdot t = 0$. Notice, that we will always multiply the value $t \cdot r$ from an unsatisfied clause with t to the right. Hence, by associativity the resulting matrix will always be zero.

If all clauses are fulfilled the value computed is t^{k+1}. However, t is idempotent in $R^{2\times 2}$ and hence is equal to t. □

Let P be the product program as in Theorem C.1. Then it is obvious that

$$\#\text{3-SAT}(\phi) = \left(\sum_{e \in \{0,1\}^m} P_\phi(e) \right)_{(1,1)},$$

the first entry in the resulting matrix, as every satisfying assignment contributes exactly one and every unsatisfied assignment zero. To compute the sum with the Cayley determinant we will use the following special matrix form.

Definition C.2. *We say a $n \times n$ matrix M is a* barber pole matrix *if it is of the form*

$$M[i,j] = \begin{cases} \alpha_i & \text{if } i = j, \\ \beta_i & \text{if } i = j + 1 \mod n, \\ 0 & \text{otherwise.} \end{cases}$$

for α_i, β_i non zero.

Notice, that for every barber pole matrix there exists only two cycle covers. Either the one where every vertex takes a self-loop or the single cycle.

Theorem C.2 ([14, Theorem 5])**.** *The value $\sum_{e \in \{0,1\}^m} P_\phi(e)$ can be computed by the Cayley determinant of a matrix of size $n \times n$ over the algebra $R^{2\times 2}$.*

Proof. Let $P = (1, (\iota_1, a_{1,0}, a_{1,1}), \ldots, (\iota_n, a_{n,0}, a_{n,1}))$. Let

$$\mathcal{I}_\ell = \{i \in [n] \mid \text{the } i\text{th instruction uses the } \ell\text{th bit of the input}\}.$$

Let $\mathcal{I}_\ell$ have the instructions $i_{\ell,1}, \ldots, i_{\ell,|\mathcal{I}_\ell|}$. Let π_0 be the identity permutation and $\pi_1(i_{\ell,\kappa}) = i_{\ell,\kappa+1 \mod |\mathcal{I}_\ell|}$ the "shifted" permutation. Notice, that this corresponds to multiple cyclic permutation, consisting of cycles of length $|\mathcal{I}_\ell|$ for $2 \leq \ell \leq n$ where the elements of the cycle are the elements in $\mathcal{I}_\ell$. They are ordered in the natural order of the instructions. Furthermore, we pad every set $\mathcal{I}_\ell$ to have size at least two.

Then we define the matrix

$$M[i,j] = \begin{cases} (-1)^{|\mathcal{I}_k|-1} a_{i,0} & \text{if } i = i_{k,1},\ j = i_{k,2} \text{ for some } k, \\ a_{i,b} & \text{otherwise, if } j = \pi_b(i), \\ 0 & \text{otherwise.} \end{cases}$$

Left to show is that $\mathsf{C\text{-}det}(M)$ is indeed computing the value of the product program. Let us look at this matrix a bit closer and assume the position $M[i_{k,1}, i_{k,2}]$ is $a_{i,0}$. We can see that this is a block barber pole matrix where the entry not on the diagonal are permuted.

Let us generate the matrix only for the set of instructions $\mathcal{I}_\ell = \{i_{\ell,1}, \ldots, i_{\ell,|\mathcal{I}_\ell|}\}$. At first we add all the entries for π_0 which are just on the diagonal. If we now look at π_1, we see that the first entry we add is at position $(1, i_{\ell,2})$ where $i_{\ell,2}$ is the index of the next instruction. We continue this until π_1 wraps around. This is clearly a cycle in the graph represented by the matrix.

It is now clear that π_1 produces a cycle for every $\mathcal{I}_\ell$. In essence it enforces that we either take all self-loops or all elements corresponding to π_1. Meaning we either multiply all values in the instructions asking x_ℓ for x_ℓ being zero or all values for the instructions where x_ℓ is one.

By this argument it is clear that one cycle cover is the value of the product program where we chosen every bit of the input and hence the value of all cycle cover is $\sum_{e \in \{0,1\}^m} P(e)$. Now it is easy to see that the actual value of $M[i_{k,1}, i_{k,2}]$ exactly cancels the sign the determinant introduces. □

As the Cayley permanent is equal to the Cayley determinant for this construction we get the following corollary.

Corollary C.1. *The value $\sum_{e \in \{0,1\}^m} P_\phi(e)$ can be computed by the Cayley permanent of a matrix of size $n \times n$ over the algebra $R^{2\times 2}$ by removing the scalar factor of $(-1)^{|\mathcal{I}_k|-1}$.*